

LCA Methodology

Numerical Approaches Towards Life Cycle Interpretation

Five Examples

Reinout Heijungs, René Kleijn

Centre of Environmental Science, Leiden University, P.O. Box 9518, NL-2300 RA Leiden, The Netherlands

Corresponding author: Dr. Reinout Heijungs; e-mail: heijungs@cml.leidenuniv.nl

DOI: <http://dx.doi.org/10.1065/lca2000.12.045>

Abstract. The ISO-standard for LCA distinguishes four phases, of which the last one, the interpretation, is the least elaborated. It can be regarded as containing procedural steps (like a completeness check) as well as numerical steps (like a sensitivity check). This paper provides five examples of techniques that can be used for the numerical steps. These are the contribution analysis, the perturbation analysis, the uncertainty analysis, the comparative analysis, and the discernibility analysis. All five techniques are described at a non-technical level with respect to basic concept, possibilities, tabular and graphical representation, restriction and warnings, and all are illustrated with a simple example.

Keywords: Comparative analysis; contribution analysis; discernibility analysis; life cycle interpretation; perturbation analysis; sensitivity analysis; statistical techniques; uncertainty analysis

Introduction

Developments within and surrounding the establishment of an ISO standard for LCA have spurred the introduction and conceptual development of a new phase of the LCA framework interpretation. In earlier texts, like SETAC's 'Code of Practice' (Consoli et al. 1993), the US guidelines (Vigon et al. 1993) and the Dutch guidelines (Heijungs et al. 1992), there was a phase of initialisation (goal definition and/or scoping), two, three or four phases in which the processing and collecting of data took place (inventory analysis and impact analysis/assessment, where the latter was sometimes divided into classification, characterisation and (e)valuation), and one optional phase in which the possibilities for product or process improvement are investigated (improvement analysis). The idea of an improvement analysis always instigated discussion: improvement was felt to be an application of LCA, and therefore not a part of LCA. Others found a way out by separating the analysis of possible improvements from the actual implementation of improvements. The Dutch guidelines, for instance, provided a numerical method (the marginal analysis) which presented a list of key issues to which small changes might lead to substantial improvements in environmental performance. However, the approach was not generally applied, perhaps partly for reasons of its lack of clarity. In the context of SETAC, the improvement analysis has never been worked out in the way that inventory analysis, impact assessment and data quality have been.

The disputed character of the improvement assessment, and the need for a systematic place for a discussion of the uncertainties of LCA results have made it evident that the phase of improvement assessment has been superseded by the phase of interpretation. According to ISO, interpretation is "a systematic technique to identify, qualify, check, and evaluate information from the results of the life cycle inventory (LCI) analysis and/or LCIA of a product system" (ISO 1999). In the draft ISO standard¹ for interpretation, there are quite a number of possibilities for interpretation. For the purpose of this paper, we may categorise these into a class of procedural approaches and numerical approaches. The procedural approaches include all types of analyses that deal with the data and results in relation to other sources of information, like expert judgements, reports on similar products, intuition, reputation of data suppliers, and so on. With the numerical approaches, we will capture those approaches that somehow deal with the data that is used during the calculations, without reference to those other sources of information. In other words, the numerical approaches explore the data in different ways. In general, LCA can be seen as a form of data reduction: thousands of numbers enter the calculation, and only a dozen or one hundred are reported as the LCA-results. This leads to a loss of information. We now define numerical approaches towards interpretation as algorithms that use and process the data in different ways, so as to produce different types of 'smart' data reduction that provide an indication of reliability, key issues, discernibility, robustness, and so on.

In this paper, we present five concrete ways to explore the data that is used in an LCA. All these explorations may take place at the level of the inventory analysis, the characterisation, the normalisation or the final weighting. The numerical approaches towards interpretation could be used as a part of each of those steps or phases, or after the whole sequence of phases. All five approaches are discussed with respect to the basic concept, the spectrum of possibilities, the most appropriate ways of tabulating or visualising the results, and the restrictions and pitfalls. The five approaches are

- contribution analysis
- perturbation analysis
- uncertainty analysis
- comparative analysis
- discernibility analysis.

¹ The draft standard has become a standard during the review period of this article. The quoted sentence has changed a bit in the final standard.

Their definition and meaning is discussed in the following five sections. The last section handles the five approaches in an encompassing discussion. Some of the techniques (e.g. the contribution analysis) are quite familiar, while others (e.g. the discernibility analysis) are entirely new. In addition to that, well-known techniques have not often been described in their full range of possibilities and along with their restrictions. Finally, a discussion of the relative strengths and weaknesses of the different techniques seems to be lacking anyway. Therefore, this paper is not primarily focussed on novel techniques, even though it describes some. Main emphasis is on systematically describing and comparing five complementary approaches.

Every approach is illustrated with a simple example. For this purpose, software has been used. The software that has been used, as well as the data that are used, are available on the Internet (<http://www.leidenuniv.nl/cml/ssp/cmllca>), so that the interested reader has the opportunity to make a closer study of the five numerical approaches described in this paper, and to explore the different options. The uncertainty analysis and the discernibility analysis use random numbers, so that the results reported in the examples cannot be reproduced exactly. The example is a hypothetical comparison of three product alternatives for producing light: the incandescent lamp, the fluorescent lamp, and the tube lamp. All data are completely fictitious, and are meant for illustration only.

1 Contribution Analysis

1.1 Basic concept

The first approach in this paper is the contribution analysis, which is also sometimes called dominance analysis or analysis of key issues. The idea of the contribution analysis is to decompose the aggregated results of inventory analysis, characterisation, normalisation or weighting into a number of constituent elements. For instance, one may wish to investigate the share of electricity production to the total carbon dioxide emission of a product life cycle. The idea of a contribution analysis is so obvious that it has been practised in many case studies, and that it is mentioned in many methodological treatises, although a clear exposition has not often been written. We mention in this respect Heijungs et al. (1992) and ISO (1999).

There can be several purposes for doing a contribution analysis. Knowing the share of a certain process or life cycle stage in a certain emission or impact category may provide opportunities for the redesign of products or processes, or for prevention strategies at a more general level. This is an application-oriented use of the contribution analysis. But there are also analysis-oriented uses. A contribution analysis points out those elements that make the highest contribution to a certain emission or impact category, and a precise knowledge of the data that correspond to those elements is therefore a prerequisite for a precise LCA result. Conversely, a rough estimate is likely to be acceptable for those elements that hardly contribute. However, it should be kept in mind that 'false negatives' due to underestimated or missing flows cannot be identified with a contribution analysis. A further

use of the contribution analysis is for testing the results against what one would intuitively expect. If the LCA of car transport is dominated by the use of the car radio, there is probably a severe error in one or more data entries.

1.2 Possibilities

The contribution analysis may be used at the level of inventory analysis, characterisation, normalisation and weighting. With higher levels of aggregation, there are more directions along which a decomposition into contributing elements may also be performed.

At the inventory analysis level, there is not much choice. One may here investigate the contributions of the various unit processes that form the life cycle. Alternatively, one may assign all unit processes to a smaller number of life cycle stages, like production of materials, production of energy, use and maintenance, and post-consumer treatment, in order to decompose the inventory results into the contributions of those life cycle stages.

At the characterisation and normalisation level, there is one other direction of decomposition. One may investigate the share of unit processes (or life cycle stages) or the share of elementary flows in a category result. And one may even combine these two directions of decomposition. Thus, one may decompose the acidification score into the contributions of electricity production, tyre production, waste incineration, and so on, or into the contributions of NO_x , SO_2 , NH_3 , and so on, or into the contribution of electricity production through NO_x , electricity production through SO_2 , tyre production through SO_2 , waste incineration through NO_x , and so on.

At the weighting level, there is even one more direction: decomposition into the contribution of impact categories. One may now investigate the share of unit processes (or life cycle stages) or the share of elementary flows or the share of impact categories in the weighted index. And one may now combine two or all three of those directions, for instance investigating the share of electricity production associated with NO_x emissions causing acidification in the weighted index.

1.3 Tabular and graphical representation

The results of a contribution analysis are the contributions that certain unit processes (or life cycle stages), elementary flows and/or impact categories make to an aggregated LCA-result. As such, they can be expressed in percentages that add up to 100. This can be visualised easily with pie charts. A table with the contributions, sorted in descending order is also insightful.

There is one important complication that most notably turns up in a pie chart: it may happen that certain contributions are negative. The three most important cases here are unit processes with negative emissions (like forestry which removes CO_2 from the air), elementary flows with negative characterisation factors (like NO which has in some scenarios a negative POCP), and unit processes that are sub-

tracted from a product life cycle, for reasons of coproduct allocation with the substitution method, leading to avoided emissions and avoided resource extractions. A stacked bar diagram, in which negative contributions are shown downwards, and in which the 100% line may be lower than the top of the upwards stack will be more easily understood by the reader than the pie chart. The tabular representation does not suffer from the problem of negative contributions, although the entire concept itself will always require explanation to the naive reader.

1.4 Restrictions and warnings

If one uses a contribution analysis to investigate the share of the unit process 'use of a refrigerator' in the impact category 'global warming', one finds a small contribution through the leakage of CFCs or HCFCs, but not a large contribution through the release of CO₂ in electricity production. That is because the unit process 'use of a refrigerator' does not release CO₂; only the unit process 'production of electricity' does so. This follows from a straightforward implementation of the concept of contribution: use of a refrigerator does not directly contribute to the CO₂ emission, but only indirectly. One might wish to include the indirect contribution by redefining the concept of the contribution analysis, but this leads to very strange results. Use of a refrigerator also indirectly implies the depreciation, and thereby the replacement and dismantling of the refrigerator itself. Therefore, inclusion of indirect contribution ultimately leads to contributions of 100%, thereby depriving the contribution analysis of any meaning.

Although a contribution analysis can take place on clusters of unit processes, the life cycle stages, it cannot take place on clusters of elementary flows, like 'heavy metals', 'chlorinated hydrocarbons' or 'pesticides'. Doing so would require a rule on how to add heavy metals: in mass units, in toxic equivalents, or in whatever way. It is thus necessary to select an elementary flow as a focal point for a contribution analysis at the inventory analysis level, to select an impact category as a focal point at the characterisation and normalisation level, and, without selection, the weighted index at the weighting level.

1.5 Example

For the product alternative fluorescent lamps, **Table 1** shows the contribution of the different processes to the atmospheric carbon dioxide emission. We see that the production of electricity is the process with the highest share (67%) in the life cycle CO₂ emission.

Table 1: Example of the results of a contribution analysis. See text

unit process	kg	%
production of electricity	0.002	67
incineration of fluorescent lamps	0.0008	27
production of fuel	0.0002	7

2 Perturbation Analysis

2.1 Basic concept

The second approach that we discuss is the perturbation analysis. This approach has been introduced as the marginal analysis by Heijungs et al. (1992) and, more formally, by Heijungs (1994). It is in some respects close to what may be referred to as a sensitivity analysis, but this term has to most people not a formal definition and meaning, and is applicable to any analysis that explores the sensitivity of a calculation result. The basic idea of the perturbation analysis is that small (marginal) perturbations of the input parameters propagate as smaller or larger deviations of the resulting output, and that a knowledge of which parameters lead to large deviations and which lead to small deviations may be useful.

There are two main purposes for a perturbation analysis. One is that it provides a checklist of those input parameters of which a small imprecision already leads to important changes in the results. Thus, it draws the attention of the researcher to those data items that should be known most precisely, whereas it also lists those data items of which even large uncertainties are unimportant, and that therefore do not deserve priority in a more detailed analysis. The second purpose of the perturbation analysis is application-oriented. Knowledge of the sensitive data items may suggest ideas for product and process improvement. If one knows that a 1% change of the electricity use of the production process leads to 4% less CO₂, a careful consideration of the electric efficiency seems natural. Conversely, if a 1% change of the amount of transport of the product leads to only 0.001% less resource use, it seems best to concentrate the improvement process to different items than the logistics details.

The extent to which the perturbation of a certain input parameter propagates into a certain output result can be interpreted as a multiplier. If an increase of 1% of an input parameter leads to an increase of 2% of an output result, the multiplier that connects those two items is said to be 2. If the output result decreases by 2%, the multiplier is said to be -2. In theory, the concept of multipliers is restricted to marginally small changes. Thus, it is not necessarily true that a change of input of 40% with a multiplier of 2 leads to a change of output of 80%. In certain situations, this will be the case, but in other cases it will certainly not. Multipliers for the perturbation analysis may be found by a complicated analytical formula (Heijungs 1994), or by numerical methods, simply calculating a result with and without a perturbed parameter, and dividing the difference in results by the difference in input parameters. In the example below, the size of the perturbation has been set to 1.001, which means that calculations were performed with a certain value and with that value increased by 0.1%.

Most multipliers will be between -1 and 1, with a concentration around 0, although values smaller than -1 and larger than 1 may occur under certain conditions (and are in fact of special interest); see under Possibilities. As a rule of thumb, one can say that multipliers of which the absolute value is higher than 0.8 and especially larger than 1 are noteworthy. For reasons that go beyond the scope of this paper, the multipliers for one selected elementary flow, impact category, or the weighted index, add to 0.

It is worth observing that the perturbation analysis does not require that parameter uncertainties be specified. It makes an analysis of the inherent sensitivity of the results for each consecutive input parameter, without paying regard to the real uncertainty of these parameters. The perturbation analysis can therefore be performed whenever LCA-results are produced; no additional data are required. When uncertainty estimates of input parameters are available, they are ignored in the perturbation analysis, and their use is discussed in the sections on uncertainty analysis and discernibility analysis. As we will see, there is a large similarity between the contribution analysis and a perturbation analysis of elementary flows. This similarity may easily lead to a misapprehension of the entire concept of perturbation analysis. Its main interest lies in pointing out the system's response to small changes of the economic flows between unit processes.

2.2 Possibilities

The perturbation analysis can be performed at the four levels that are distinguished throughout this paper: inventory analysis, characterisation, normalisation and weighting. However, with increasing input data being needed as the level of aggregation grows higher, more input data items are required, and more input data items can hence be perturbed. At the inventory analysis level, one selects an elementary flow, and can choose to perturb the intermediate flows of products, materials and energy (which we will call the economic flows) or to perturb the elementary flow at the successive unit processes. Interestingly enough, it is the recursive structure of economic flows (coal production needs electricity, and electricity production needs coal), which makes certain that multipliers are larger than 1 or smaller than -1. Also interesting is that the perturbation of the input parameters that correspond to the elementary flows of the unit processes gives a result that is very similar to the result of the contribution analysis. Therefore, one could argue that the perturbation analysis encapsulates the contribution analysis, thereby obviating the need for the latter one. From a practical perspective, however, things are different. The perturbation analysis is much slower and more difficult to comprehend than the contribution analysis. Moreover, the unit processes are perturbed one by one, and their resulting multipliers are listed one by one. A grouping into life cycle stages seems incompatible with the idea of the perturbation analysis, while it is obvious for the contribution analysis.

At the characterisation level, the characterisation factors can also be perturbed; at the normalisation level, the normalisation factors can be perturbed, and at the weighting level, the weighting factors can be perturbed. In principle, the perturbation analysis can also be applied to allocation factors.

2.3 Tabular and graphical representation

As the multipliers higher than 0.8 and smaller than -0.8 are of particular interest, the most obvious tabular presentation is to sort the absolute values of the multipliers in descending order, with an optional cut-off for small multipliers, say between -0.2 and 0.2. A graphical representation is no more insightful than a tabular one.

2.4 Restrictions and warnings

An important restriction on the use of the perturbation analysis is the time it takes. If one uses the numerical approximation, a product life cycle of 100 unit processes, with each unit process connected with 6 economic flows, the perturbation analysis for one elementary flow requires 600 calculation procedures. When one calculation takes one second, which seems to be optimistic for certain LCA software implementations, this means 10 minutes computing time. Computation time will increase slightly at higher levels of aggregation, as the most time-intensive step in the algorithm is the inventory analysis. The analytic formulas for the perturbation are not likely to be less time consuming, as they imply a bunch of time consuming matrix manipulations, again with 600 repetitions.

2.5 Example

Table 2 shows part of the results of a perturbation analysis of the economic flows with respect to atmospheric carbon dioxide emission for the product alternative fluorescent lamps. We see that producing 1% more light in the use process yields 0.999% less CO₂, and that producing 1% more electricity in the electricity production process yields 0.733% less CO₂. Using 1% less electricity in the use process results in 0.731% less CO₂.

Table 2: Example of the results of a perturbation analysis. See text

unit process	economic flow	multiplier
use of fluorescent lamps	fluorescent lamp light	-0.999
production of electricity	electricity	-0.733
use of fluorescent lamps	electricity	0.731
use of fluorescent lamps	disposed fluorescent lamp	0.266
incineration of disposed fluorescent lamps	disposed fluorescent lamp	-0.266
production of electricity	fuel	0.067
production of fuel	fuel	-0.067

3 Uncertainty Analysis

3.1 Basic concept

There are many input parameters that are known to be uncertain. Several data sources yield different production or emission characteristics, or different levels of environmental impact. The systematic study of the propagation of input uncertainties into output uncertainties will take place in the uncertainty analysis. For several reasons, the topic will be treated on the basis of Monte Carlo simulations (see, e.g. Morgan & Henrion 1990).

Suppose that we have a system with input parameters that are specified in the form of a probability density function. For instance, for each parameter one knows that the value is distributed according to a Gaussian or normal distribution, with a certain mean value and a certain standard deviation. In one Monte Carlo simulation, one obtains one outcome of the system by drawing every parameter value from its distribution. In this way, one finds a certain result,

say for the aggregated CO₂ emission. In another Monte Carlo run, one obtains another outcome by drawing again from the specified distributions. This will lead to a new value of the CO₂ emission. If this procedure of drawing and calculating is repeated a large number (say 1000) of times, a probability density function for the aggregated CO₂ emission will have been constructed. It can be characterised by a certain mean, a standard deviation, and possibly other statistical indices (like the median). Possibly, it follows more or less a Gaussian distribution, so that mean and standard deviation suffice to give a complete description.

The purpose of a large number of Monte Carlo simulations is to provide an understanding of the uncertainty of the LCA results. A specific output result, like '120 kg CO₂' then receives some sort of indication of significance (although not significance in the statistical sense), like '120 kg CO₂ with a standard deviation of 10 kg'. A comparison of product alternatives on the basis of merely results without uncertainties is harder to believe than if uncertainties are included. Statements on uncertainty intervals may in principle improve to this situation; see, however, the section on discernibility analysis for a more complete discussion.

3.2 Possibilities

All input data can include information on uncertainties. This applies to unit process data, to characterisation factors, to normalisation factors, and to weighting factors. That is, one may specify all input data in the form of probability density functions. In practice, this reduces to only a few different options: a normal distribution with specified mean and standard deviation, and a uniform distribution with specified smallest and largest values are the two most frequent. Huijbregts (1998) also uses triangular and lognormal distributions for certain parameters. Sometimes, one may wish to specify two or three discrete values with specific probabilities, for instance a car with one, two, three or four passengers with probabilities 0.5, 0.3, 0.1 and 0.1. In principle, one may also indicate uncertainty estimates of allocation factors, and even to allocation principles, for instance substitution method, mass allocation and economic allocation with probabilities 0.4, 0.3 and 0.3. In that case, one should interpret the term 'probability' as something like the 'degree of belief'.

3.3 Tabular and graphical representation

The simulation results may be presented in different forms. The most basic form is to calculate mean and standard deviation for every output result. Observe that the mean may differ from the baseline result without uncertainty estimates. The difference may be due to chance, but it may also be that the distribution of the output result is asymmetric, so that mean and mode differ, even for extremely large numbers of simulations. The reason for the asymmetry is explained in the Appendix.

One may also choose to produce more statistics for every output item: lowest and highest value, mean, mode and median, standard deviation and skewness. One may even produce test statistics for normality (the Kolmogorov-Smirnov

test), for difference with another product alternative (the t test), and so on. A problem is, of course, that more statistics means more pages of output, and that interpretation should provide a help rather than a bunch of pages filled with statistical information. Another alternative is to tabulate all realisations for every Monte Carlo trial. The most important graphical aid is in that case a histogram of the Monte Carlo results, which gives a quick indication of how these values are distributed. Aspects like symmetry and unimodality can then easily be detected visually.

3.4 Restrictions and warnings

As with the perturbation analysis, the uncertainty analysis requires quite some computing time. It is said in literature (Morgan & Henrion 1990) that 10,000 runs yield in general reliable results. This means calculating 10,000 times the inventory analysis and possibly the subsequent impact assessment. With the aforementioned estimate of 1 second per inventory calculation, some 3 hours would be needed. This suggests that a researcher makes some quick calculations with 100 runs, and that the real analysis with 10,000 runs is something to have run overnight.

An uncertainty analysis presumes that uncertainty parameters are available for all input parameters, like unit process data, characterisation factors, normalisation factors, weighting factors, allocation details, and so on. These uncertainty parameters should be quantitative information, like standard deviations or other parameters that specify a probability distribution. Qualitative labels, like the pedigree matrix of Weidema and Wesnæs (1996), do not suffice to carry out an uncertainty analysis in this sense. Unfortunately, there is at present not much available with respect to quantitative uncertainty parameters. Most available databases with unit processes and equivalency factors do not contain standard deviations, and an LCA-practitioner is often glad to get point estimates of data at all, and will not pursue getting interval estimates or other measures of dispersion.

3.5 Example

Table 3 presents the results of an uncertainty analysis with 1000 runs of the atmospheric emission of carbon dioxide of the product alternative incandescent lamp. We see that the result without uncertainty analysis (indicated as the baseline result) coincides well with the mean of the Monte Carlo runs. This suggests a quite symmetrical distribution, which indeed is confirmed by a graphical inspection. The coefficient of variation is defined as the ratio between the stand-

Table 3: Example of the results of an uncertainty analysis. See text

parameter	value	unit
baseline	0.024	kg
mean	0.024	kg
standard deviation	0.002	kg
variation	8	%
lowest	0.018	kg
highest	0.031	kg

ard deviation and the mean; a value of 8% suggests a reasonably certain result. This is also suggested by the extreme values, which lie about 0.06 kg on both sides of the mean value, again suggesting a fairly symmetrical distribution.

Fig. 1 presents the frequency distribution of the 1000 different values for the carbon dioxide emission that were obtained by the uncertainty analysis. We see that the distribution is fairly symmetric, resembles a normal distribution, and has a maximum that is quite close to the baseline result.




Fig. 1: Histogram (with bin size 0.005) of the individual Monte Carlo results of an uncertainty analysis. See text

4 Comparative Analysis

4.1 Basic concept

The comparative analysis is nothing more than a systematic place to list the LCA results for different product alternatives simultaneously.

4.2 Possibilities

The comparative analysis can take place at all four levels, i.e. inventory analysis, characterisation, normalisation and weighting. The most interesting feature that we mention here is that different scales can be used to display the results. First, one may show the absolute values, e.g. the kg CO₂, the kg SO₂, etc., all in their own units and scales. Alternatively, one can put the smallest or the largest result for each elementary flow, impact category or weighted index to 1. For instance, if one puts the largest elementary flow to 1, one easily sees how much better each alternative is for each elementary flow. Or, one may put all elementary flows for one product alternative to 1, thereby declaring that alternative as the reference product.

4.3 Tabular and graphical representation

All the possibilities that are mentioned find an easy place in a tabular form with a column for every product alternative and a row for every elementary flow or impact category.

A natural graphical presentation is through a bar chart for a certain result, where each bar represents the score of one product alternative. A logarithmic scale is often difficult to interpret, but one should be cautious with any product alternative that strongly dominates the others, thereby suggesting that the difference between the other alternatives is small.

4.4 Restrictions and warnings

A comparative analysis is seductively simple. It is dangerous, because it may easily induce one to make claims without a proper analysis of the robustness of these claims with respect to the influence of uncertainties.

4.5 Example

Table 4 shows an example of the results of a comparative analysis at the inventory level for the three product alternatives. We have chosen to put the lowest intervention to 1 for every elementary flow, so that we can easily see how much worse a certain product alternative is. For instance, the tube lamp is superior for CO₂, as it beats the fluorescent lamp by a factor of 2 and the incandescent lamp by a factor of more than 16.

Table 4: Example of the results of a comparative analysis. See text

elementary flow	incandescent lamp	fluorescent lamp	tube lamp
CO ₂ to air	16.2	2.0	1
SO ₂ to air	27.7	2.8	1
copper to soil	1	1.1	1.1
sand	2.5	1	1.4
copper ore	1	6.0	5.7
crude oil	27.7	2.8	1

5 Discernibility Analysis

5.1 Basic concept

It is said that an important goal of LCA is the comparison of product alternatives, but a comparison need not assume the form of comparison of point estimates and/or interval estimates. What matters, in fact, is the ranking of product alternatives in a statistically sound way, like judgements of the form 'product A is significantly better than product B'. Statistically significant then follows the usual interpretation of 'there is a 95% chance that product A is better than product B', at least if the significance level is put at a conventional 0.05 level. In other words, we seek to test if product A is statistically discernible from product B.

The idea of the discernibility analysis stems from the desire to combine the comparative analysis and the uncertainty analysis. One Monte Carlo realisation is used to calculate the results for all product alternatives simultaneously. Notice that this means that the discernibility analysis is only applicable when uncertainty estimates are available. Huijbregts (1998) gives an example for two product alternatives. He proposes the comparison index (CI) as the ratio between the score for the two alternatives, and shows frequency distributions for the CI. If a significant part (e.g. 95%) of the frequency distribution is on one side of the 1, the two alternatives are said to have significantly different scores. As a comparison index is the ratio between the scores of two products, use of it is restricted to comparing two product alternatives only.

The discernibility analysis is a generalisation for the case of two or more product alternatives. To understand the idea, it

is necessary to note that, although Huijbregts' approach calculates the difference (or rather the ratio) of the scores in one Monte Carlo run, and although the full probability distribution is constructed, the final judgement discards most of this information. The only thing that counts in one run is, whether the score for the first product alternative is higher or lower than the score for the other one. The approach effectively comes down to counting the number of times that the first product alternative has a higher score and the number of times that the second product alternative has a higher score. Let us indicate the first event with $n(A>B)$ and the second one with $n(B>A)$. The decision criterion for discernibility of product alternatives A and B with respect to the selected item (emission, impact category, weighted index) is that either of the n s dominates the other in a statistically significant way. A usual significance level is 0.95, so if there are n runs, $n(A>B)$ should be at least $0.95n$ if we are allowed to declare that product alternative A has a significantly higher score than B. And conversely, $n(B>A)$ should be at least $0.95n$ if we are allowed to declare the opposite. If $n(A>B)$ is $0.95n$ or less and $n(B>A)$ is $0.95n$ or less as well, we may say that we cannot reject the null hypothesis of indiscernibility of the two product alternatives. Obviously, two almost indiscernible product alternatives will yield about $0.50n$. And in normal situations, the probability of a tie (A and B having exact the same score) is vanishingly small, so that $n(A>B)+n(B>A)=n$.

If the direction (smaller or larger) and not the distance between the scores for the product alternatives is recorded, it leads to an important reduction of information. It is exactly the reduction that creates the possibility of analysing the general case of more than two product alternatives. See the section on representation below for an example.

5.2 Possibilities

Like any Monte Carlo-based analysis, the uncertainties of the input data and a number of runs are required. The discernibility analysis can take place for elementary flows, for impact categories, and for the weighted index. The major option for displaying results, is that one can choose between listing counts (like $n(A>B) = 9,163$ for 10,000 runs), listing fractions (like $p(A>B) = 0.9163$), listing percentages (like $p(A>B) = 91.63\%$), and listing significance tests (like $p(A>B) = \text{n.s.}$, where n.s. means not significant at the 95%-level).

5.3 Tabular and graphical representation

A convenient form of representing the results of a discernibility analysis is as in Table 5. In this table, one reads that the score for product alternative B is lower than that for A in 91% of the Monte Carlo runs, that it is lower than the score for product alternative C in 43% of the runs, and that the score for product alternative C is lower than that for A in 97% of the runs. Hence, with a significance level of 95%, product C can be said to have a significantly lower score than product A. Furthermore, product B can be said to have a quite but not significantly lower score than product A. Finally, products B and C are barely discernible with respect to the selected indicator. Of course, one could restrict the table to the right upper triangle alone, as the other part is redundant.

Table 5: Illustration of the results of a discernibility analysis. Here, one would conclude that product C has a lower score than product A in 97% of the simulations, so that the two products are highly discernible, whereas product C is almost indiscernible from product B

	A	B	C
A	—	0.09	0.03
B	0.91	—	0.43
C	0.97	0.57	—

One could visualise the ranks of the different products on a line interval, indicating significance intervals. It is questionable if this would add much to the tabular presentation.

5.4 Restrictions and warnings

A point to notice is that the discernibility analysis ignores the distance between the scores of product alternatives; it only uses a smaller-larger dichotomy. It measures the probability that a specific product alternative has a lower (or higher) score than the other alternatives. Two alternatives can be very close in numerical value (say, 12.2 and 12.3 kg CO₂) and still be statistically discernible or vice versa. The difference is due to the degree of uncertainty in the point estimates. In that sense, the discernibility analysis fits in the set of non-parametric statistical tests, like the sign test, the Kruskal-Wallis test, and Kendall's coefficient of concordance (see, e.g. Siegel 1956). In a certain sense, the discernibility analysis is a complement to the comparative analysis: the former only yields information with respect to discernibility, while the latter only states the point estimates of the scores.

Because the discernibility analysis is a special form of an uncertainty analysis, the problems that are associated with this latter type of analysis also apply here. A discernibility analysis is time-consuming, and requires the specification of many uncertainty parameters.

5.5 Example

Table 6 shows the results of a discernibility analysis at the characterisation level for the impact category ecotoxicity. We see that the tube lamp has a lower score than the fluorescent lamp in all cases, which suggests that the scores are fully discernible and that tube lamps are with 100% certainty superior to fluorescent lamps with respect to ecotoxicity. The tube lamp has a lower score than the incandescent lamp in about 40% of the cases. This suggests that the incandescent lamp is better, but not significantly better. A similar argument can be made for the discernibility of fluorescent lamps and incandescent lamps.

Table 6: Example of the results of a discernibility analysis. See text

	incandescent lamp	fluorescent lamp	tube lamp
incandescent lamp	—	66	60.4
fluorescent lamp	34	—	0
tube lamp	39.6	100	—

6 Discussion

This paper has given an overview of five different ways to explore the data of an LCA. The five approaches can be used under different conditions, that are summarised in

Table 7. Having performed a point estimate of the elementary flows, the impact category results, the normalisation results, or the weighted index, one can always do a contribution analysis and a perturbation analysis, and we think that it is wise to include them in every LCA that has the ambition to transcend from a quick scan. Most software for LCA includes the possibility for a contribution analysis, but unfortunately does not include a perturbation analysis. We hope that practicality and practice will change in this respect.

Table 7: Overview of five numerical approaches towards life cycle interpretation that are discussed in this paper in relation to their applicability along the dimensions for one product/several products and with/without uncertainty estimates

	one product alternative	more than one product alternative
without uncertainty estimates	contribution analysis perturbation analysis	comparative analysis
with uncertainty estimates	uncertainty analysis	discernibility analysis

We have discussed that uncertainty estimates, required for an uncertainty analysis and a discernibility analysis, is most often lacking. Carrying out Monte Carlo simulations therefore implies either using partial uncertainty information (e.g. for a selection of unit processes and/or characterisation factors), or adding subjective uncertainty estimates (like putting the standard deviation to 10% where no information is available). Both options are obviously second-rate solutions. On the one hand, the majority of LCA-software programmes cannot deal with uncertainty estimates of input parameters because these data have typically not been available anyhow. On the other hand, the absence of computational devices for handling such information has lowered the priority for the collection of these data. We think it is important that both designers of LCA-software and developers of LCA-databases for inventory analysis and impact assessment show the courage to escape from this trap. In this respect, we refer to Burmaster & Anderson (1994) for a survey of the principles of good practice for carrying out Monte Carlo simulations.

If the purpose of the LCA is to rank product alternatives in a decision-context, the right-hand column in Table 7 becomes of interest. If the purpose is merely an analysis of a certain product life cycle, with possible applications in the form of obtaining recommendations, the other column is important.

A combination of techniques may be appropriate in certain cases. One could, for instance apply a perturbation analysis to find out for which data obtaining uncertainty intervals will have the highest priority. A subsequent uncertainty analysis and/or discernibility analysis will then be more powerful.

The subject of interpretation has barely been addressed in literature. This paper provides five numerical techniques, and the ISO standard and certain other reports provide a number of more procedural techniques. We think that there may be many more numerical approaches for interpretation. We also think that the five approaches that have been discussed are in their infancy, which means that they may be developed further, and that their usefulness and realm of application is not yet completely clear. Although we have benefited much from the ISO text on interpretation, it should be clear that the label 'standard' is somewhat premature, in the sense that it is diffi-

cult to speak of a standard at a time that there is merely a procedural framework, and no experience with the application. More research efforts in this field are needed, especially if one recognises that the input data of a typical LCA consists of thousands of numbers, and that statistical techniques for data exploration have been developed to a high degree of sophistication in other fields of interest. In this respect, we just mention the enormous toolbox of multivariate methods, like principal component analysis, factor analysis and cluster analysis (see, e.g. Johnson & Wichern 1992).

References

Burmaster DE, Anderson PD (1994): Principles of good practice for the use of Monte Carlo techniques in human health and ecological risk assessments. *Risk Analysis* 14 (4) 477-481

Consoli F, Allen D, Boustead I, Fava J, Franklin W, Jensen AA, de Oude N, Parrish R, Perriman R, Postlethwaite D, Quay B, Seguin J, Vigon B (1993): Guidelines for life-cycle assessment: a 'Code of Practice'. Edition 1. SETAC, Brussels/Pensacola

Heijungs R, Guinée JB, Huppes G, Lankreijer RM, Udo de Haes HA, Wegener Sleeswijk A, Ansems AMM, Eggels PG, van Duin R, de Goede HP (1992): Environmental life cycle assessment of products. Backgrounds – October 1992. CML, Leiden

Heijungs R (1994): A generic method for the identification of options for cleaner products. *Ecological Economics* 10 (1) 69-81

ISO (1999): Environmental management – Life cycle assessment – Life cycle interpretation. FDIS 14043. ISO, Geneva

Huijbregts MAJ (1998): Application of uncertainty and variability in LCA. Part II. Dealing with parameter uncertainty due to choices in life cycle assessment. *International Journal of LCA* 3 (6) 343-351

Johnson RA, Wichern DW (1992): Applied multivariate statistical analysis. Third edition. Prentice-Hall, Inc., Englewood Cliffs

Morgan MG, Henrion M (1990): A guide to dealing with uncertainty in quantitative risk and policy analysis. Cambridge University Press, New York

Siegel S (1956): Non-parametric statistics for the behavioral sciences. McGraw-Hill Book Company, Inc., New York

Vigon BW, Tolle DA, Cornaby BW, Latham HC, Harrison CL, Boguski TL, Hunt RG, Sellers JD (1993): Life-cycle assessment. Inventory guidelines and principles. EPA, Cincinnati

Weidema BP, Wesnæs MS (1996): Data quality management for life cycle inventories. An example of using data quality indicators. *Journal of Cleaner Production* 4 (3-4) 167-174

Received: May 15th, 2000

Accepted: November 14th, 2000

Online-First: December 8th, 2000

Appendix

Symmetric uncertainty distributions of input parameters may easily lead to asymmetric uncertainty distributions of results. This appendix demonstrates for a very simple system how this works.

Suppose that a product system consists of two unit processes: production of a product and production of electricity. To make 1 product, one needs 10 MJ electricity, and to make 1 MJ electricity one emits 10 kg CO₂. There will therefore be a baseline emission of 100 kg CO₂. Now suppose that we replace the point estimate of 10 MJ electricity per product by a symmetric interval estimate that ranges between 5 and 15 MJ electricity per product. It follows that this leads to an emission of CO₂ that ranges between 50 kg and 200 kg. This interval is obviously not symmetrically centred around the baseline result of 100 kg. Observe that this argument nowhere makes assumptions with respect to the exact distribution: it applies to uniform and Gaussian and all other symmetric distributions.